



Homeland  
Security

# Measuring DoS Impact

Jelena Mirkovic, University of Delaware

Sonia Fahmy, Purdue University

Peter Reiher, UCLA

Roshan Thomas, SPARTA

Alefiya Hussain, SPARTA

Steven Schwab, SPARTA

Calvin Ko, SPARTA



# Goals

---

- ❖ To measure DoS defense effectiveness we must first measure DoS attack impact
  - Effectiveness measure then expresses how quickly and completely this impact is removed by the defense

# Existing Approaches

---

- ❖ Percentage of good packets arriving at the victim
  - Includes retransmissions, does not capture delay or jitter which are detrimental to some traffic
- ❖ Percentage of attack packets dropped
  - Worst: says nothing about legitimate traffic
- ❖ Duration and packet loss
  - Less is better but how much is low enough?
- ❖ Throughput
  - More is better but how much is high enough?
- ❖ Goodput
  - Works only for TCP

# Main Challenges

---

- ❖ Different applications have different QoS requirements
  - One-way delay, request/response delay, jitter, throughput, packet loss, overall duration
- ❖ These requirements are very subjective
  - Even same person may find different tolerance thresholds depending on the task it performs, her expectations and the service provider's reputation
- ❖ We must define thresholds that are
  - Measurable
  - Define success or failure for majority of people
  - Acceptable to research and commercial communities

# Application QoS Requirements

---


- ❖ Use CAIDA's traffic statistics to determine popular applications
- ❖ Use QoS research to define QoS requirements and appropriate thresholds
- ❖ Luckily, we can borrow many results from 3GPP initiative
  - 3GPP is a “collaboration agreement which brings together a number of telecommunications standards bodies” from all over the world, in an effort to “produce globally applicable Technical Specifications and Technical Reports for a 3rd Generation Mobile System”
  - Advantage: large standard bodies already agreed with these QoS specifications

# Application QoS Requirements

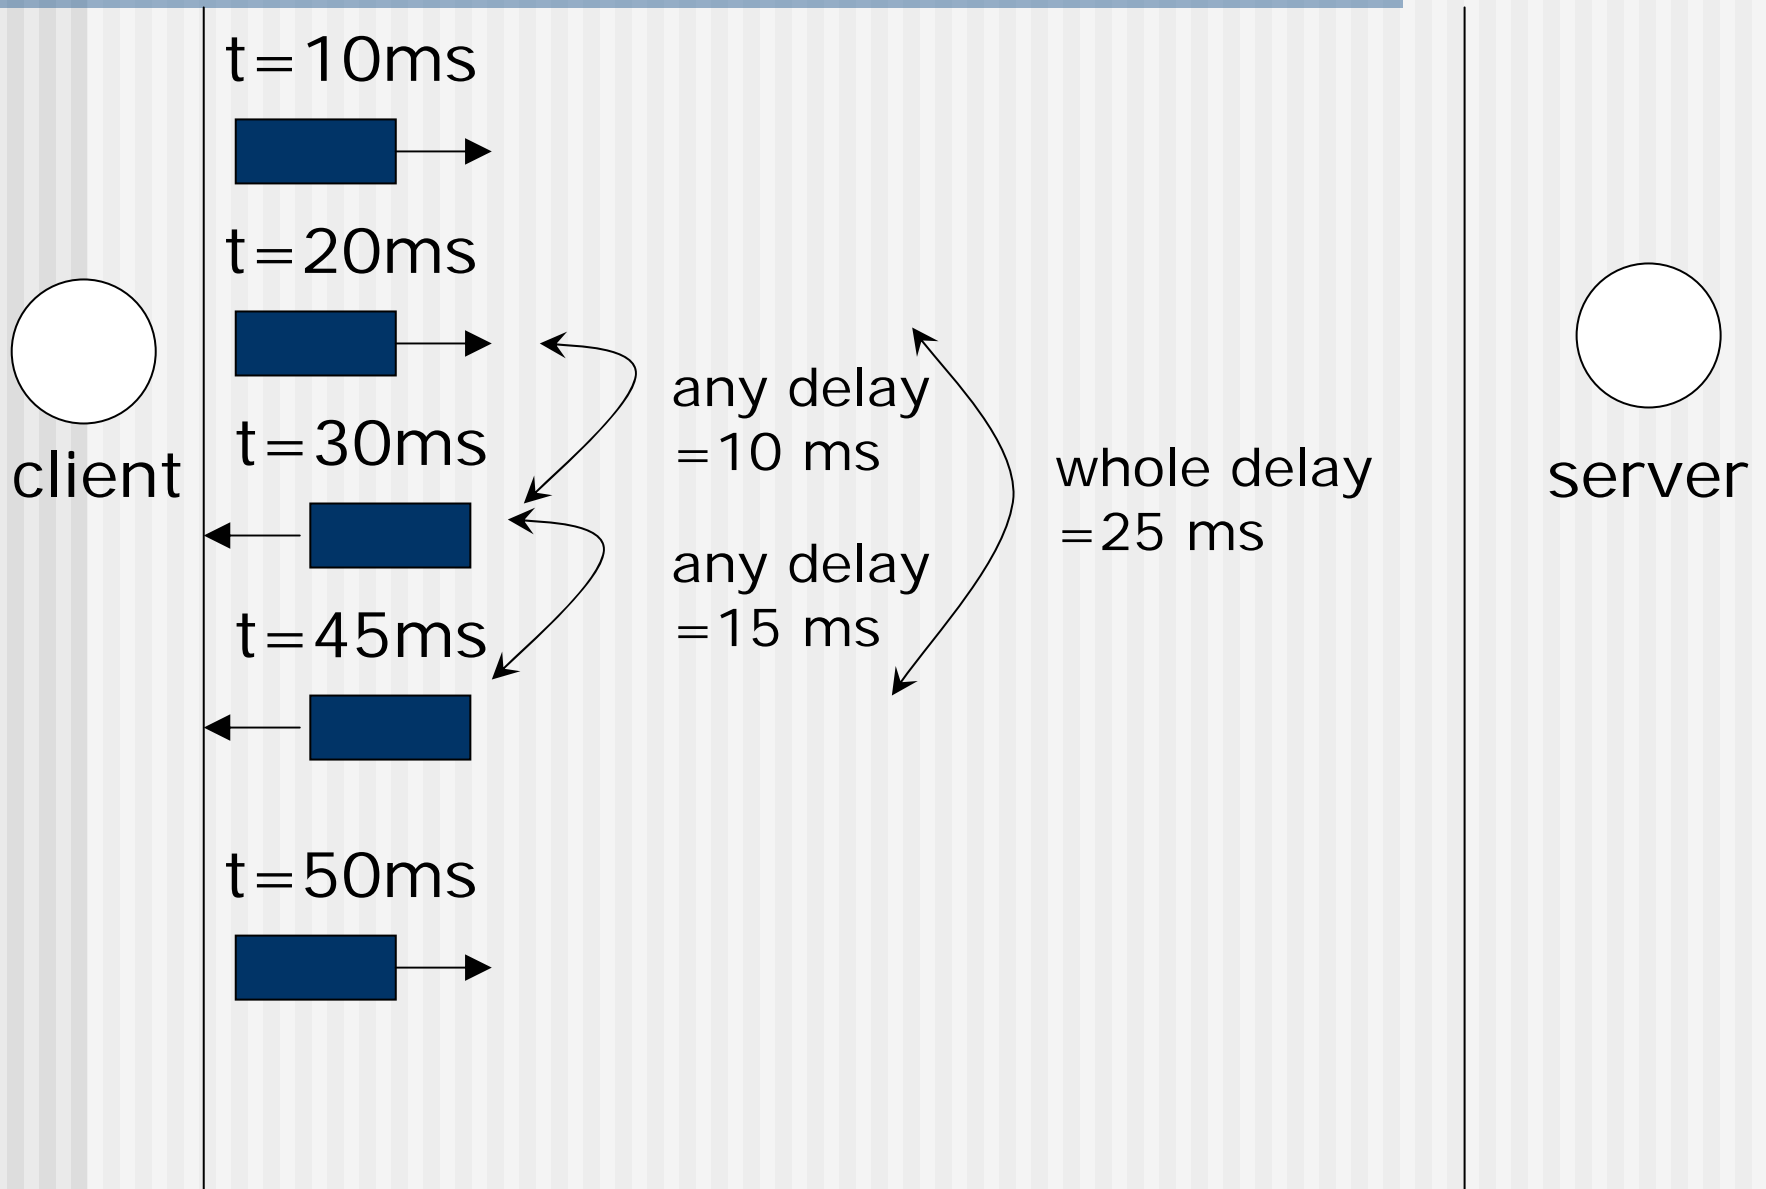
Category	One-way delay	Req/rep delay	Loss	Duration	Jitter
E-mail (server-server)		whole < 4 h			
NNTP		whole < 4 h			
Chat	< 30 s				
Web		any < 4 s		< 60 s	
FTP, file sharing		any < 10 s		< 300%	
FPS games	< 150 ms		< 3%		
RTS games	< 500 ms				
Telnet		any < 250 ms			
E-mail (user-server)		any < 4 s		< 300%	
DNS		whole < 4 s			
ICMP		whole < 4 s			
Audio, convers.	< 150 ms	whole < 4 s	< 3%		< 50 ms
Audio, voice mesg.	< 2 s	whole < 4 s	< 3%		< 50 ms
Audio, stream	< 10 s	whole < 4 s	< 1%		< 50 ms
Videophone	< 150 ms	whole < 4 s	< 3%		
Video, stream	< 10 s	whole < 4 s	< 1%		

# One-way Delay

---



# Request/Response Delay



# Defining Transactions

---

- ❖ If an hour-long Telnet connection does not meet its QoS requirements just before the end does this mean that it has failed?
  - No, user had good service for almost an hour
  - A part of this connection failed
- ❖ We define **transaction** as some high-level task a user wants to perform
  - Downloading a file
  - Having a VoIP conversation
  - Browsing one Web page
- ❖ We measure transaction success and failure

# Defining Transactions

Application	Transaction
E-mail (server-server)	TCP flow
NNTP	TCP flow and inactive > 4 s
Chat	TCP flow and inactive > 4 s
Web	TCP flow and inactive > 4 s
FTP, file-sharing	TCP flow and inactive > 4 s
Games	UDP flow
Telnet	TCP flow and inactive > 4 s
E-mail (user-server)	TCP flow and inactive > 4 s
DNS	Request/response
ICMP	Request/response
Audio and video	TCP flow and its UDP flow

# Measuring DoS Impact

---

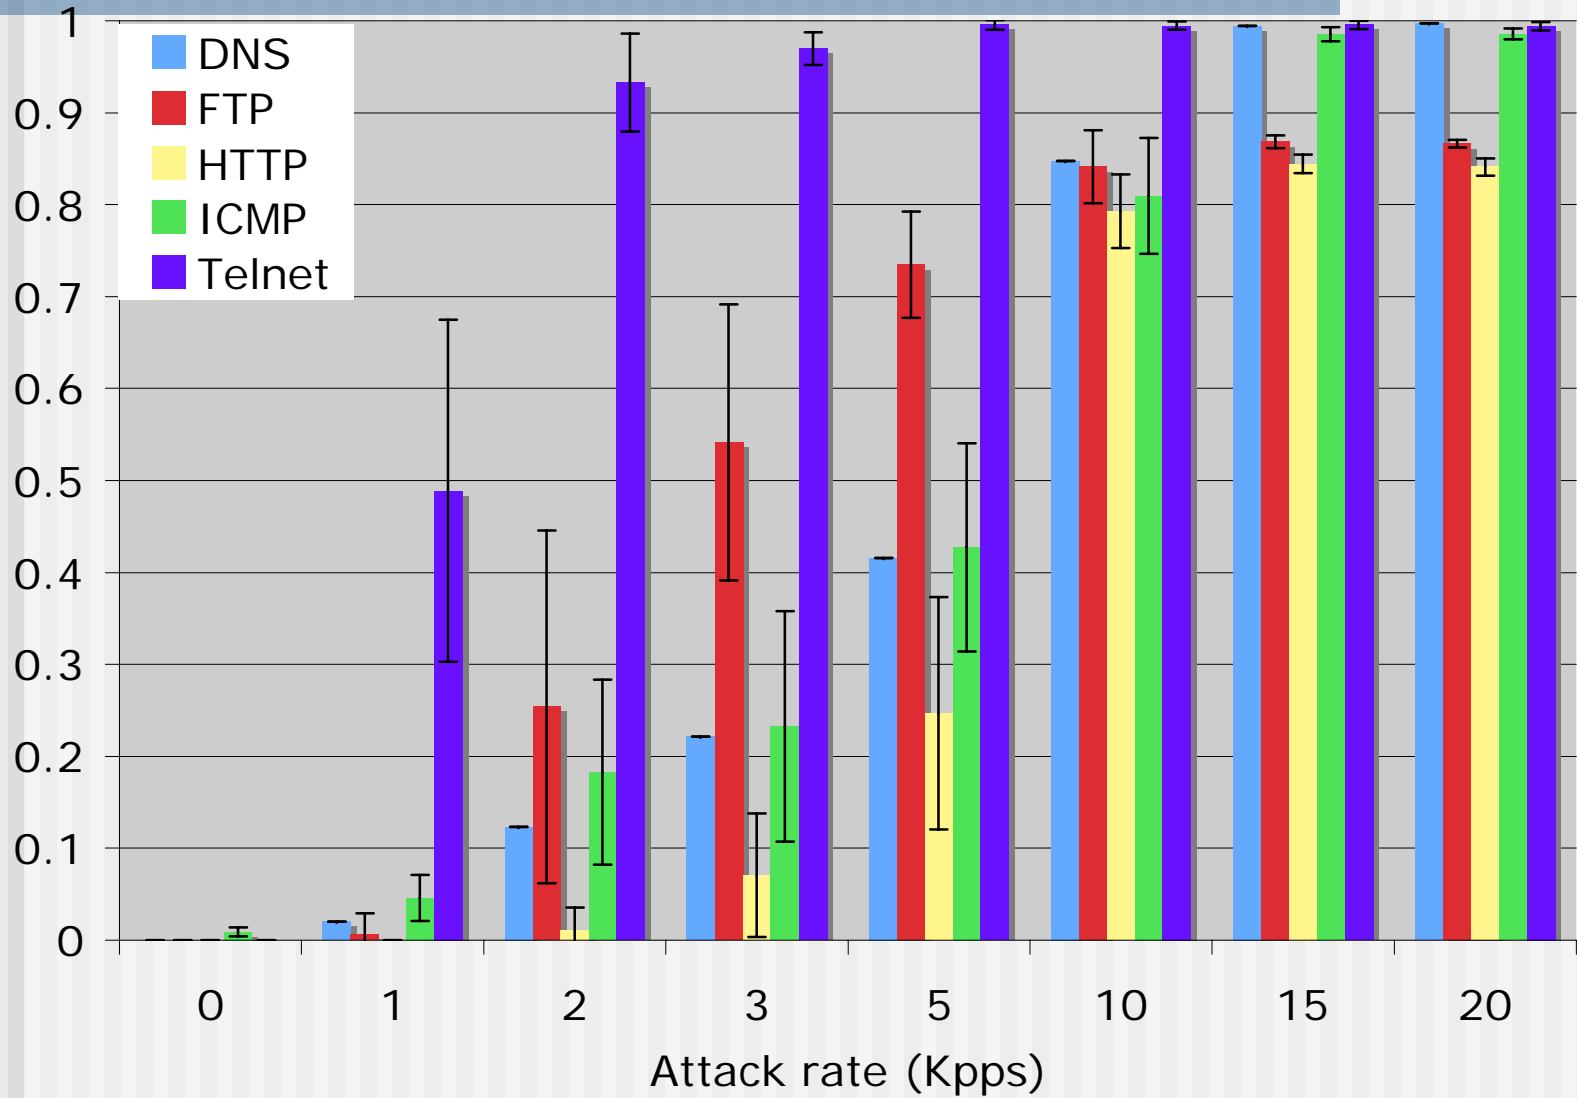
- ❖ We calculate percentage of failed transactions *pft* for each application category
- ❖ **DoS-hist** measure: histogram of *pft* values across categories or destination ports
- ❖ **DoS-degree** measure: aggregates *pft* values across categories using weights into a single number
- ❖ **Timed measures**: We can measure transaction failure in time to capture time dimension of a defense

# Illustration

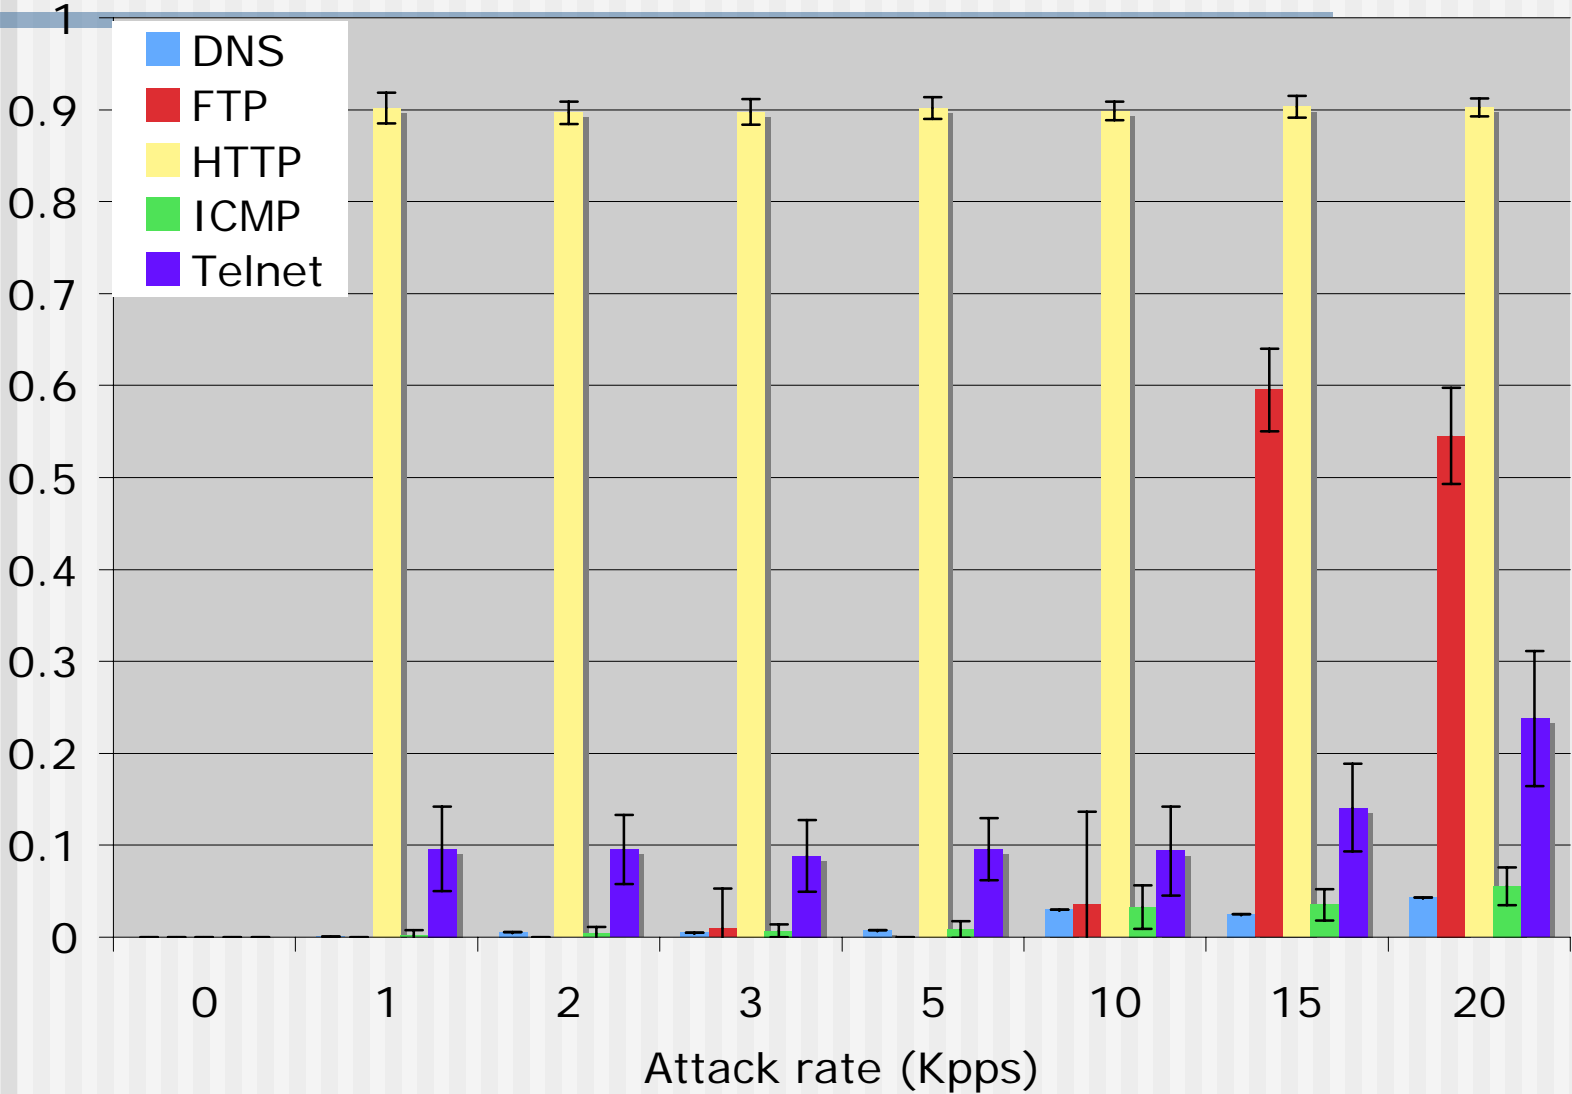
---

- ❖ We created a simple DETER experiment
  - One legitimate client, one attacker, one target with a bottleneck link
- ❖ Traffic:
  - Telnet with Pareto-distributed duration and volume and exponential arrivals
  - FTP with Pareto-distributed file sizes and exponential arrivals
  - HTTP with Pareto-distributed file sizes and exponential arrivals
  - DNS with exponential arrivals
  - ICMP with exponential arrivals


# UDP Flood



# TCP SYN Flood on Web service




### Goodput




UDP flood

### Goodput

### TCP SYN flood



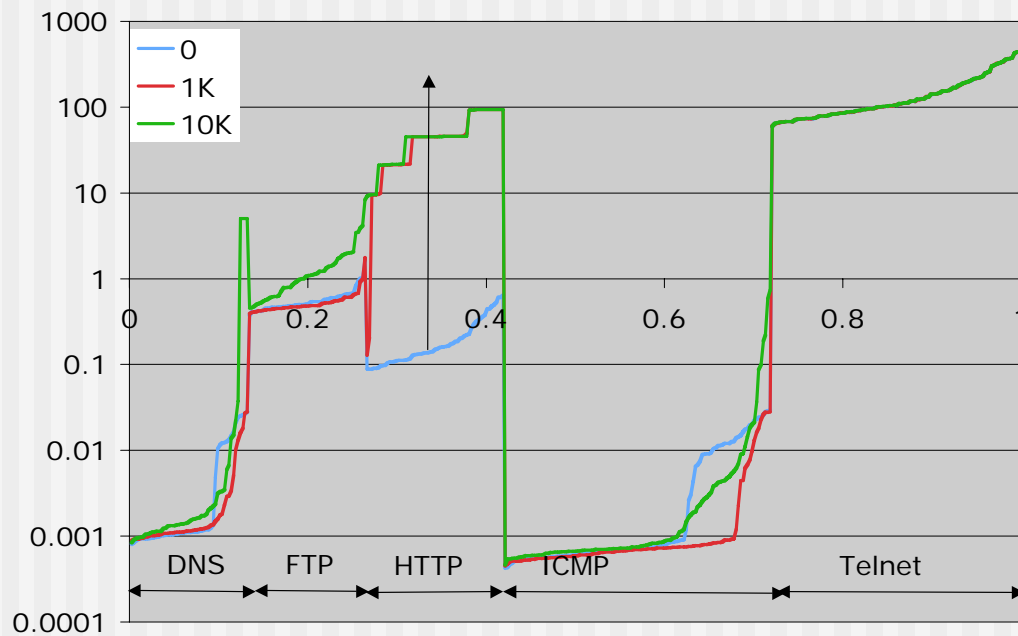
Duration vs attack strength



UDP flood

Percentage transactions

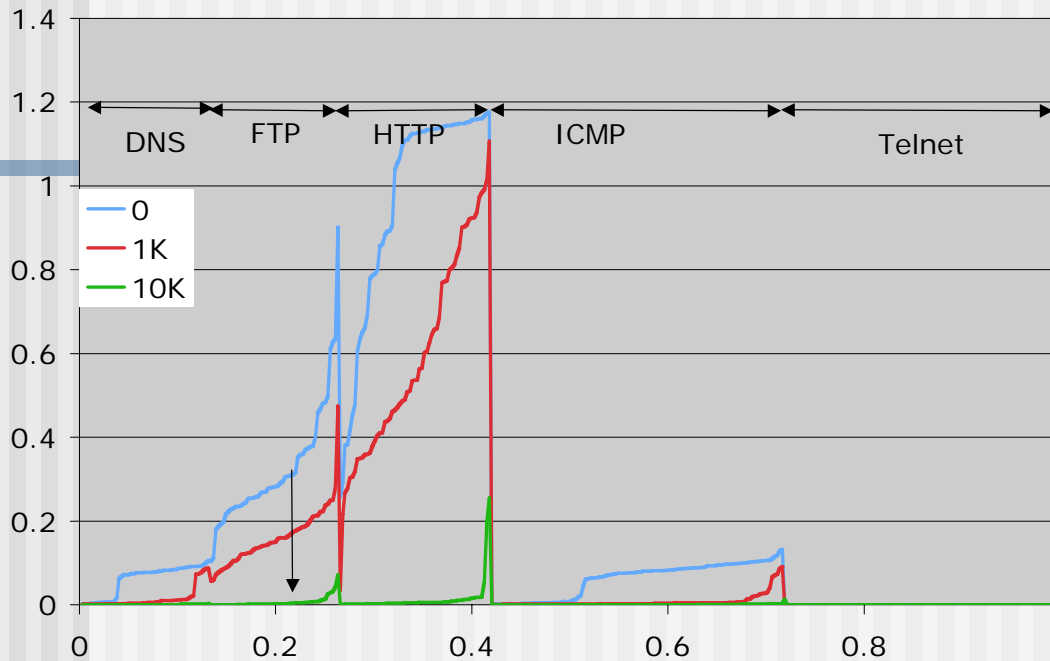
Duration vs attack strength



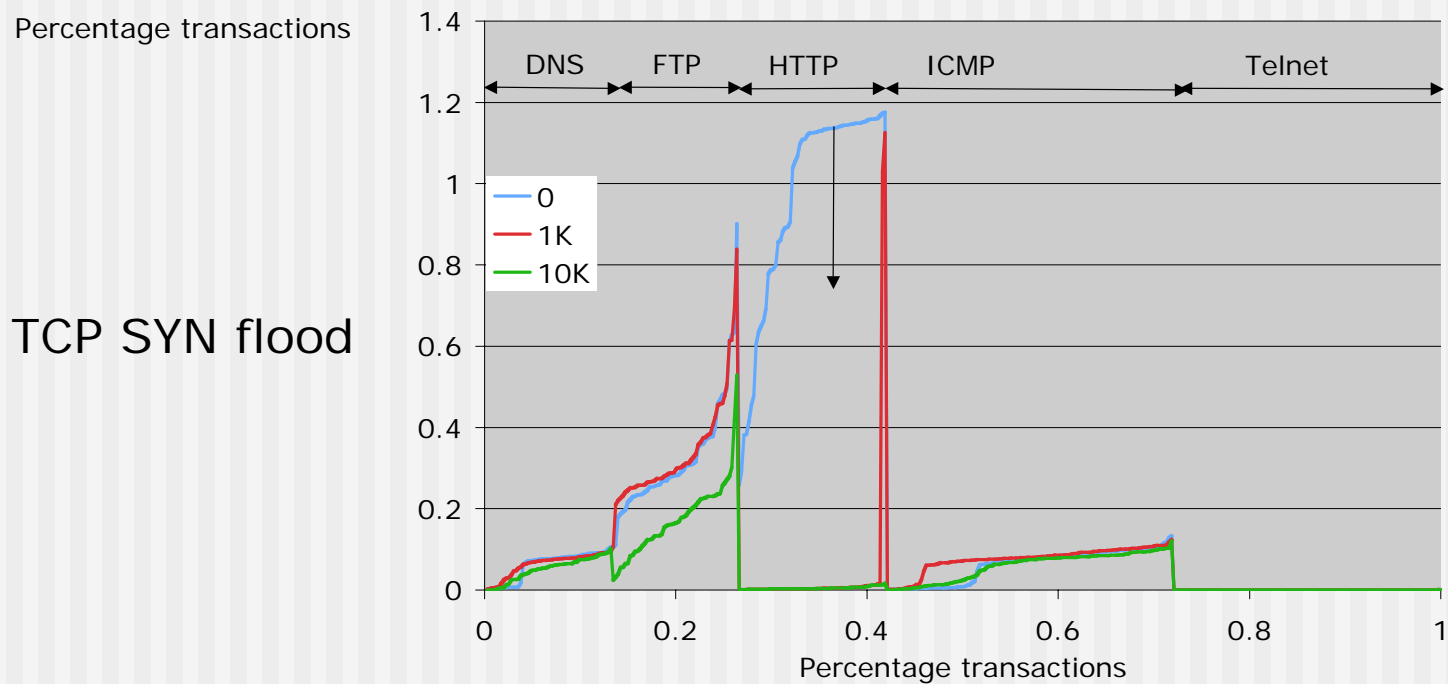
TCP SYN flood

Percentage transactions

Throughput vs attack strength



Throughput vs attack strength



# Conclusions

---

- ❖ Many open issues remain:
  - ❖ Using tcpdump at high packet rates
  - ❖ How to aggregate measures into a single number
  - ❖ Are these thresholds realistic?
- ❖ We have just scratched the surface, much work remains to be done
- ❖ DoS impact measure must be defined to accurately measure effectiveness of defenses

Questions? Comments?  
sunshine@cis.udel.edu

# HTTP goodput

